Computers and Operations Research 89 (2018) 17–30

Contents lists available at ScienceDirect

Computers and Operations Research

journal homepage: www.elsevier.com/locate/cor

A linear programming based algorithm to solve a class of optimization problems with a multi-linear objective function and affine constraints

Hadi Charkhgard[a,*], Martin Savelsbergh[b], Masoud Talebian[c]

[a] Department of Industrial and Management Systems Engineering, University of South Florida, USA
[b] School of Industrial and Systems Engineering, Georgia Institute of Technology, USA
[c] Graduate School of Management and Economics, Sharif University of Technology, Iran

ARTICLE INFO

Article history:
Received 10 November 2016
Revised 19 July 2017
Accepted 31 July 2017
Available online 2 August 2017

Keywords:
Pareto optimal solutions
Convex programming
Multi-linear objective function
Linear programming
Polynomial-time algorithm

ABSTRACT

We present a linear programming based algorithm for a class of optimization problems with a multi-linear objective function and affine constraints. This class of optimization problems has only one objective function, but it can also be viewed as a class of multi-objective optimization problems by decomposing its objective function. The proposed algorithm exploits this idea and solves this class of optimization problems from the viewpoint of multi-objective optimization. The algorithm computes an optimal solution when the number of variables in the multi-linear objective function is two, and an approximate solution when the number of variables is greater than two. A computational study demonstrates that when available computing time is limited the algorithm significantly outperforms well-known convex programming solvers IPOPT and CVXOPT, in terms of both efficiency and solution quality. The optimization problems in this class can be reformulated as second-order cone programs, and, therefore, also be solved by second-order cone programming solvers. This is highly effective for small and medium size instances, but we demonstrate that for large size instances with two variables in the multi-linear objective function the proposed algorithm outperforms a (commercial) second-order cone programming solver.

1. Introduction

In this paper, we study optimization problems of the form,

$$\begin{aligned}\max\ & \prod_{i=1}^{p} y_i \\ \text{s.t. } & \boldsymbol{y} = D\boldsymbol{x} + \boldsymbol{d} \\ & A\boldsymbol{x} \le \boldsymbol{b} \\ & \boldsymbol{x}, \boldsymbol{y} \ge \boldsymbol{0}, \quad \boldsymbol{x} \in \mathbb{R}^n, \quad \boldsymbol{y} \in \mathbb{R}^p,\end{aligned}$$

with D a $p \times n$ matrix, $\boldsymbol{d}$ a p-vector, A an $m \times n$ matrix, and $\boldsymbol{b}$ an m-vector. Such optimization problems often arise in game theory settings where the $\boldsymbol{x}$ variables represent players' actions and the $\boldsymbol{y}$ variables represent players' utilities. Examples include computing the Nash solution to a bargaining problem (Nash, 1950, 1953) and computing an equilibrium of a linear Fisher or a Kelly capacity allocation market (Chakrabarty et al., 2006; Eisenberg and Gale, 1959; Jain and Vazirani, 2007; Vazirani, 2012a).

We will refer to such an optimization problem as a *positive multi-linear program with affine constraints* (PMP-A). A PMP-A can both be seen as an extension of a linear programming problem and as a special case of a geometric programming problem. We refer to the set $\mathcal{X} := \{\boldsymbol{x} \in \mathbb{R}^n : A\boldsymbol{x} \le \boldsymbol{b},\ \boldsymbol{x} \ge \boldsymbol{0}\}$ as *the feasible set in the decision space* and to the set $\mathcal{Y} := \{\boldsymbol{y} \in \mathbb{R}^p : \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{y} = D\boldsymbol{x} + \boldsymbol{d},\ \boldsymbol{y} \ge \boldsymbol{0}\}$ as *the feasible set in the payoff space*. We assume that $\mathcal{X}$ is bounded (which implies that $\mathcal{Y}$ is compact) and that the optimal objective value of the problem is strictly positive, i.e., there exists a $\boldsymbol{y} \in \mathcal{Y}$ such that $\boldsymbol{y} > \boldsymbol{0}$. We usually refer to $\boldsymbol{x} \in \mathcal{X}$ as a *feasible solution* and to $\boldsymbol{y} \in \mathcal{Y}$ as a *feasible point* ($\boldsymbol{y}$ is the image of $\boldsymbol{x}$ in the payoff space). We note that, at first glance, the problem appears to be closely related to multiplicative programs (for more information on multiplicative programs, see for instance Gao et al. (2006); Ryoo and Sahinidis (2003); Shao and Ehrgott (2014,2016). Specifically, by changing the objective function of a PMP-A from *max* to *min* a multiplicative program is obtained. However, for this to be a reformulation, it is necessary to require that the variables in the objective function take on only non-positive values, which violates the definition of a multiplicative program. Moreover, it is known that a multiplicative program is NP-hard (Shao and Ehrgott, 2016), but a PMP-A can be solved in polynomial time. Furthermore, for multiplicative programs, it is known that there exists an optimal

* Corresponding author.
E-mail addresses: hcharkhgard@usf.edu, h.charkhgard@gmail.com (H. Charkhgard).

http://dx.doi.org/10.1016/j.cor.2017.07.015
0305-0548/

solution that is an extreme point of polyhedron $\mathcal{Y}$ (Shao and Ehrgott, 2016), which can be exploited to develop custom-built algorithms. Unfortunately, this property does not hold for a PMP-A.

It can be shown that a PMP-A has a *unique* optimal point (Nash, 1950). We denote an *optimal solution* and the *optimal point* of a PMP-A by $\boldsymbol{x}^*$ and $\boldsymbol{y}^*$, respectively. Finally, we assume that $\boldsymbol{x}^*$ and $\boldsymbol{y}^*$ are rational. Moreover, it is easy to see that $\boldsymbol{x}^*$ is an *efficient (or Pareto optimal) solution*, i.e., a solution in which it is impossible to improve the value of a player's payoff without a deterioration in the value of at least one other player's payoff, and, consequently, $\boldsymbol{y}^*$ is a *nondominated point* (the image of an efficient solution in the payoff space is called a nondominated point) (Ehrgott and Gandibleux, 2007; Eusébio et al., 2014).

A convex programming solver, for example one that uses an interior point method, can find an optimal solution to a PMP-A in polynomial time (Grötschel et al., 1988). Note that because we have assumed that the optimal value is positive, it is possible to use $\sum_{i=1}^{p} \log y_i$ as the objective function. However, because convex programming solvers are significantly slower than (commercial) linear programming solvers, it may be possible to develop a linear programming based algorithm that performs better in practice. That is exactly what we do.

The main contribution of our research is the development of a linear programming based (LP-based) algorithm for solving instances of PMP-A. When the number of variables in the multi-linear objective function is two (i.e., $p = 2$), the algorithm computes an optimal solution in polynomial time, and when the number of variables is greater than two the algorithm computes an approximate solution (but with a quality guarantee). A computational study demonstrates that when available computing time is limited our algorithm significantly outperforms the well-known convex programming solvers IPOPT and CVXOPT. A PMP-A can be reformulated as a second-order cone program (SOCP), and, therefore, solved by any SOCP solver. The SOCP approach works well and outperforms the LP-based algorithm on small to medium size instances. For large size instances with $p = 2$, however, the LP-based algorithm still dominates.

Vazirani (2012b) was the first to present an LP-based algorithm for solving a PMP-A with a bilinear objective function (i.e., $p = 2$), the Binary Search Algorithm (BSA). However, BSA cannot be easily extended to handle a multi-linear objective function with more than two variables (i.e., $p > 2$). Both BSA and our algorithm work in the payoff space, but in different ways. BSA searches the payoff space for a facet of $\mathcal{Y}$ containing the optimal point (note that in two dimensional space, facets are line segments) and, once found, computes the optimal point. Our algorithm searches the payoff space by iteratively removing sections of $\mathcal{Y}$ that are guaranteed not to contain the optimal point. In the process, a lower and upper bound on the optimal objective value are continuously updated until they have converged to the same value (when $p = 2$).

We note that PMP-As do not only arise in game theory. At the end of the paper, we briefly mention applications of PMPs-A in other fields of study, including geometry, statistical estimation, approximation, and polynomial programming.

The remainder of the paper is organized as follows. In Section 2, we introduce some well-known applications of PMPs-A in game theory. In Section 3, we detail the logic of our proposed algorithm. In Section 4, we report on the results of a computational study. In Section 5, we introduce applications of PMPs-A in other fields of study. Finally, in Section 6, we give some concluding remarks.

2. Positive multi-linear programs with affine constraints arising in game theory

We present three well-known game theoretic settings that give rise to PMP-As. We refer the interested readers to Chakrabarty et al. (2006); Eisenberg and Gale (1959); Jain and Vazirani (2007); Vazirani (2012a) for further information and more details.

2.1. Bargaining problems

A *bargaining problem* is a cooperative game in which all players agree to create a grand coalition, instead of competing with each other, to get a higher payoff (Serrano, 2005). To be able to create a grand coalition, the agreement of all players is necessary. Therefore, a critical question to be answered is: *What should the payoff of each player be in a grand coalition?* One of the solutions to the (symmetric) bargaining problem was proposed by Nash and is now known as the Nash bargaining solution (Nash, 1950, 1953).

We start by explaining the Nash bargaining solution in the case of two players. We denote the expected utility values of the players by $\boldsymbol{y} = (y_1, y_2)$. Let $\mathcal{Y}$ be the 2-dimensional feasible set in the payoff space containing all possible expected utility values of the players. We assume that $\mathcal{Y}$ is compact and convex, and that it is given to us by an oracle in the form of a set of affine constraints. Let $\mathcal{Y}_N$ be the nondominated frontier (i.e., the set of nondominated points) of $\mathcal{Y}$. Let the disagreement point (or status quo), $\boldsymbol{q} := (q_1, q_2)$, in the payoff space represent the payoffs that the players will receive if they do not create the coalition.

Two classical axioms imposing restrictions on a solution to a pure bargaining problem are

- *Individual Rationality:* None of the players accepts a payoff lower than the one which is guaranteed to him under disagreement, i.e., $y_1 \geq q_1$ and $y_2 \geq q_2$.
- *Pareto Optimality:* The solution must be such that the payoff for one player cannot be increased without decreasing the payoff of the other player.

Let $\mathcal{Y}^* := \{\boldsymbol{y} \in \mathcal{Y} : y_1 \geq q_1, y_2 \geq q_2\}$ and $\mathcal{Y}_N^* := \{\boldsymbol{y} \in \mathcal{Y}_N : y_1 \geq q_1, y_2 \geq q_2\}$. To satisfy the classical axioms, a bargaining solution $\boldsymbol{y}^*$ must be in $\mathcal{Y}_N^*$. However, in general, $\mathcal{Y}_N^*$ still contains an infinite number of points. Nash introduced three additional axioms:

- *Symmetry:* If $\mathcal{Y}^*$ is symmetric, i.e., for any vector $(\boldsymbol{y}, \boldsymbol{y}') \in \mathcal{Y}^*$, the vector $(\boldsymbol{y}', \boldsymbol{y})$ is also in $\mathcal{Y}^*$, then in a bargaining solution we must have $y_1^* = y_2^*$.
- *Linear Invariance:* Let $\boldsymbol{y}^*$ be a solution to a bargaining game G. Moreover, let $\hat{G}$ be a bargaining game obtained from G by an order-preserving linear transformation T of one player's utility function. The solution $\hat{\boldsymbol{u}}^*$ to the bargaining game $\hat{G}$ has to be the image of $\boldsymbol{y}^*$ under T, i.e., $\hat{\boldsymbol{u}}^* = T\boldsymbol{y}^*$.
- *Independence of Irrelevant Alternatives:* Let $\boldsymbol{q}$ be the disagreement point and $\boldsymbol{y}^*$ be a solution to the bargaining game G. Moreover, let $\hat{G}$ be a bargaining game that is obtained from G by restricting $\mathcal{Y}$ to $\hat{\mathcal{Y}}$, i.e., $\hat{\mathcal{Y}} \subset \mathcal{Y}$. If $\boldsymbol{q} \in \hat{\mathcal{Y}}$ and $\boldsymbol{y}^* \in \hat{\mathcal{Y}}$, then $\boldsymbol{y}^*$ is the solution of $\hat{G}$.

Nash proved that under the above five axioms, the optimal solution $\boldsymbol{y}^* = (\boldsymbol{y}_1^*, \boldsymbol{y}_2^*)$ to the bargaining problem is the unique point satisfying

$$\boldsymbol{y}^* \in \arg\max\{(y_1 - q_1)(y_2 - q_2) : \boldsymbol{y} \in \mathcal{Y},\ y_1 \geq q_1,\ y_2 \geq q_2\}.$$

Nash's result can be extended to pure bargaining problems with p players ($p > 2$) straightforwardly. Let $\mathcal{Y}$ be the p-dimensional feasible set in the payoff space, then for a disagreement point $\boldsymbol{q} = (q_1, \cdots, q_p)$, the optimal solution $\boldsymbol{y}^* = (y_1^*, \cdots, y_p^*)$ to the bargaining problem is a point satisfying

$$\boldsymbol{y}^* \in \arg\max\left\{\prod_{i=1}^{p}(y_i - q_i) : \boldsymbol{y} \in \mathcal{Y},\ y_i \geq q_i\ \forall i \in \{1, \cdots, p\}\right\}.$$

Kalai (1977) established that the Nash bargaining solution can be extended even further to *nonsymmetric* bargaining games. He